

## CURRICULUM VITAE

### Song Jiang

---

Performance and Architecture Lab  
Computer and Computational Sciences Division  
Los Alamos National Laboratory  
Los Alamos, NM 87545

E-mail: [sjiang@lanl.gov](mailto:sjiang@lanl.gov)  
Phone: (505) 606-0308  
Fax: (505) 667-1126  
<http://www.c3.lanl.gov/~sjiang>

### Research Interests

Operating systems, File and storage systems, Fault tolerance for high-end computing, and Distributed systems.

### Education

**Ph.D. in Computer Science**, College of William and Mary, June 2004.

Thesis Topic: *Efficient Management of Large Working Sets in Virtual Memory and File Buffers*.

Adviser: Dr. Xiaodong Zhang

**M.S. in Computer Science**, University of Science and Technology of China, May 1996.

**B.S. in Computer Science**, University of Science and Technology of China, May 1993.

### Teaching and Research Experiences

**Postdoctoral Research Associate** Los Alamos National Laboratory, since 6/2004.

**Research Intern** Microsoft Research, Redmond, WA, summer 2003.

**Research Assistant**, College of William and Mary, 5/2000 — 6/2004.

**Teaching Assistant**, College of William and Mary, 8/1999 — 5/2000.

**Lecturer**, University of Science and Technology of China, 6/1996 — 7/1999.

### Awards

**Steven K. Park Graduate Research Award** for “*original and significant work on developing the LIRS caching algorithm for virtual memory and I/O systems*”, Computer Science Department, College of William and Mary in Virginia, 2003.

## Research Projects

### 1. Incremental Checkpoint/Restart for Large-Scale Parallel Systems (see Papers [5,18])

*Building a checkpoint/restart system in Linux clusters for next generation of fault tolerant large-scale computing systems. The system provides essential functionality for transparent, highly responsive, and efficient fault tolerance based on incremental/full checkpointing at kernel level.*

### 2. Exploiting Dual Locality to Improve Hard Drive Performance (See Paper [1])

*Building a new infrastructure in memory management to maximize the sequential data accesses on disk by incorporating (1) the missing spatial locality regarding disk data locations in buffer cache management policies; and, (2) the missing temporal locality regarding history sequential block accesses in disk prefetching policies.*

### 3. Low-overhead and General-Purpose Replacement Algorithms (See Paper [2, 3, 14])

*The LIRS replacement algorithm and Clock-Pro VM replacement policy proposed in this project are being actively discussed in the Linux community. The code development progress for their inclusion in the Linux kernel is being timely reported at <http://wiki.linux-mm.org/ClockProApproximation>*

### 4. Exploiting Hierarchical Locality Caching in Multi-level Buffer Cache Hierarchy (see Papers [7, 10, 19, 21])

*The ULC multi-level caching protocol proposed in this project has been implemented and evaluated by researchers at UIUC and IBM Canada. For details see their paper on Sigmetrics 2005 titled as “Empirical evaluation of Multi-level Buffer Cache Collaboration for Storage System”.*

### 5. Thrashing Prevention and Protection and their Implementations in Linux Kernels (see Papers [13, 16, 17])

*The token-ordered thrashing protection algorithm proposed in this project has been officially adopted in the Linux kernels starting in version 2.6.10 (see <http://lxr.linux.no/source/mm/thrash.c>).*

### 6. Efficient Search in Peer-to-Peer Overlay Networks (see Papers [6, 8-12, 15, 20])

*Proposing two search algorithms called LightFlood and FloodTrail for a reliable and low traffic file search in unstructured P2P networks. They also provide low responsive times and high tolerance to transient peer populations.*

## Publications

### 1 Representative Publications

1. **Song Jiang**, Xiaoning Ding, Feng Chen, Enhua Tang, and Xiaodong Zhang, “DULO: An Effective Buffer Cache Management Scheme to Exploit Both Temporal and Spatial Locality” *Proceedings of the 4th USENIX Conference on File and Storage Technologies (FAST’05)*, San Francisco, CA, December, 2005.
2. **Song Jiang**, Feng Chen, and Xiaodong Zhang, “CLOCK-Pro: An Effective Improvement of the CLOCK Replacement”, *Proceedings of the Annual USENIX Technical Conference (USENIX’05)*, Anaheim, CA, April, 2005.

3. **Song Jiang** and Xiaodong Zhang, “LIRS: An Efficient Low Inter-reference Recency Set Replacement Policy to Improve Buffer Cache Performance”, *Proceedings of ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems (SIGMETRICS’02)*. June, 2002, Marina Del Rey, CA.

## 2 Other Conference Publications

4. Feng Chen, **Song Jiang**, Xiaodong Zhang, “SmartSaver: Turning Flash Memory into a Disk Energy Saver for Mobile Computers”. to appear in *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED’06)*, Tegernsee, Germany, October 2006.
5. **Song Jiang**, Kei Davis, Fabrizio Petrini, Xiaoning Ding, and Xiaodong Zhang, “COCA: a Locality-aware Cooperative Cache Management Protocol to Improve Network File System Performance”, to appear in *Proceedings of the 26th International Conference on Distributed Computing Systems (ICDCS’06)*, Lisboa, Portugal, July, 2006.
6. Xiaoning Ding, Dimitrios Nikolopoulos, **Song Jiang**, and Xiaodong Zhang, “MESA: Integrated Static and Run-Time Cache Management to Avoid Conflicts” to appear in *Proceedings of IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS’06)*, Austin, TX, March, 2006.
7. Roberto Gioiosa, José Carlos Sancho, **Song Jiang**, Fabrizio Petrini, Kei Davis, “Transparent Incremental Checkpointing at Kernel Level: a Foundation for Fault Tolerance for Parallel Computers”, *Proceedings of IEEE/ACM International Conference for High Performance Computing, Networking and Storage (SC’05)*, Seattle, WA, November, 2005.
8. Lei Guo, **Song Jiang**, Li Xiao, and Xiaodong Zhang, “Exploiting Content Localities for Efficient Search in P2P Systems”, *Proceedings of 18th International Symposium on Distributed Computing (DISC’04)*, Amsterdam, Netherlands, October, 2004.
9. **Song Jiang**, Xiaodong Zhang, “ULC: a file block placement and replacement protocol to effectively exploit hierarchical locality in multi-level buffer caches” *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS’04)*, Tokyo, Japan, March, 2004.
10. Lei Guo, Songqing Chen, Shansi Ren, Xin Chen, and **Song Jiang**, “PROP: a Scalable and Reliable P2P Media Streaming Proxy System” *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS’04)*, Tokyo, Japan, March, 2004.
11. Shansi Ren, Lei Guo, **Song Jiang**, and Xiaodong Zhang, “Low Overhead Topology-Aware P2P Systems”, *Proceedings of 18th International Parallel and Distributed Processing Symposium (IPDPS’04)*, Santa Fe, New Mexico, April, 2004.
12. **Song Jiang**, Xiaodong Zhang, “Efficient I/O Caching in Data Grid and Cluster Management” *Proceedings of IEEE International Conference on Cluster Computing (Cluster’03)*, December, 2003, Hong Kong, China.
13. **Song Jiang**, Lei Guo, and Xiaodong Zhang, “LightFlood: an Efficient Flooding Scheme for File Search in Unstructured Peer-to-Peer Systems”, *Proceedings of 2003 International Conference on Parallel Processing (ICPP’03)*, October, 2003, Kaohsiung, Taiwan.
14. **Song Jiang** and Xiaodong Zhang, “FloodTrail: an Efficient File Search Technique in Unstructured Peer-to-Peer Systems” *Proceedings of IEEE Globecom’03 (Globecom’03)*, December, 2003, San Francisco, CA.

15. **Song Jiang** and Xiaodong Zhang, “Adaptive Page Replacement to Protect Thrashing in Linux”, *Proceedings of the 5th USENIX Annual Linux Showcase and Conference (ALS’01)*, November, 2001, Oakland, CA.

### 3 Journal Publications

16. **Song Jiang** and Xiaodong Zhang, “Making LRU Friendly to Weak Locality Workloads: A Novel Replacement Algorithm to Improve File Buffer Cache Performance”, *IEEE Transactions on Computers*, Vol. 54, No. 8, August 2005, pp. 939-952.
17. Lei Guo, **Song Jiang**, Li Xiao, and Xiaodong Zhang, “Fast and Low Cost P2P Searching by Exploiting Localities in Peer Community and Individual Peers”, *Journal of Parallel and Distributed Computing*, Vol. 65, Issue 6, 2005, pp. 729-742.
18. **Song Jiang** and Xiaodong Zhang, “Token-Ordered LRU: An Efficient Page Replacement Policy and Implementation in Linux Systems”, *Performance Evaluation*, Volume 60, Issues 1-4, May 2005, pp. 5-29.
19. **Song Jiang** and Xiaodong Zhang, “TPF: a System Thrashing Protection Facility”, *Software - Practice & Experience*, Vol. 32, Issue 3, 2002, pp. 295-318.

### 4 Workshop Publications

20. Jose Carlos Sancho, Fabrizio Petrini, Kei Davis, Roberto Gioiosa, and **Song Jiang**, “Current Practice and a Direction Forward in Checkpoint/Restart Implementations for Fault Tolerance”, *First Workshop on System Management Tools for Large-Scale Parallel Systems* (in conjunction with IPDPS 2005), April, 2005.

### 5 Papers in Submission

21. **Song Jiang**, Kei Davis, and Xiaodong Zhang, “Coordinated Multi-Level Buffer Cache Management with Consistent Access Locality Quantification”, revised version under the second round review by *IEEE Transactions on Computers*.
22. **Song Jiang**, Haodong Wang, Lei Guo, and Xiaodong Zhang, “On Peer-to-Peer Flooding for Both Minimal Redundant Messages and Maximum Searching Scope”, revised version under the second round review by *IEEE Transactions on Parallel and Distributed Systems*.

### Professional Activities

1. Program Committee Member for ICPP’05, ICPP’06.
2. Panelist for NSF Proposal review, 2005.
3. Referee for Journals:  
*IEEE Transactions on Parallel and Distributed Systems*, *IEEE Transactions on Computers*, *Journal of Performance Evaluation*.

4. Referee for Conferences:

*USENIX'06, IPDPS'05, ICDCS'05, HPDC'05, INFOCOM'04, ICDCS'04, IPDPS'04, ICDCS'03, ICDCS'02, SC'02, Performance'02.*